

the EXTENSION

A Technical Supplement to Control Network

© 2004 Contemporary Control Systems, Inc.

Introduction To Virtual LANs

INTRODUCTION

A local area network (LAN) is a private network usually confined to one plant. Virtual LANs (VLANs) allow a single physical LAN to be partitioned into several smaller logical LANs. VLANs limit the broadcast domain, improve security and performance and are ideal for separating industrial automation systems from information technology systems.

Structured Wiring

One of the advantages cited for migrating to Industrial Ethernet from fieldbus technology is found in the comment "our plant is already wired for Ethernet. I do not need to run specialized wiring since twisted-pair wiring is already in place." This could be true since Ethernet cabling installations typically follow structured wiring standards such as TIA/EIA-568-A *Commercial Building Telecommunications Cabling Standard*. Following the standard, end stations at each work area would be wired to patch panels in a wiring closet. These patch panels would also connect to repeating hubs or switching hubs mounted in the wiring closet (Fig. 1). The cross connection between end stations and hub ports are made with short patch cords. Out of each wiring closet is a single connection to a cascaded hub located in an equipment room. All wiring closet feeds go to the equipment room. It is possible to have more than one equipment room, but it is the intent of the standard to limit the number of levels of hierarchy. It is quite possible that the plant floor is wired in a similar fashion and, in this way, all stations within the plant share the same LAN.

Sharing the same LAN may not always be a good idea. LANs are typically maintained by the information technology (IT) department that has become increasingly more interested in a secure network than maximizing up-time. Disconnecting a user suspected of having a faulty station by removing a patch cord is typically done and is treated as an inconvenience to the user. However, the same action done to a device on an industrial control system could be disastrous. Therefore, it has been suggested to have two LANs—one for IT and one for industrial automation systems. This would certainly remove the security concerns of the IT department, but segregating the physical wiring may not be possible nor convenient.

There is another reason to separate the information technology LAN and the industrial automation system

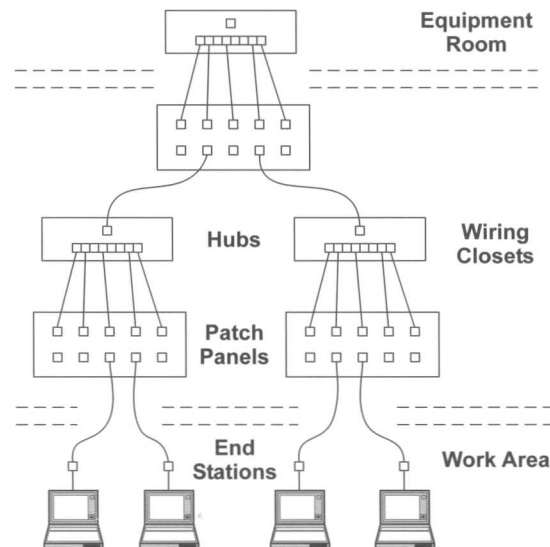


Figure 1. Structured wiring creates a hierarchy of hubs.

LANs. A LAN is considered a single broadcast domain. This means that broadcast messages (messages destined to all stations) will be sent to every station on the LAN. This is usually true for multicast messages (messages destined to many, but not all stations). If the exact location of stations that are to receive a multicast message is not known, then all stations will receive the message. Industrial automation protocols frequently use the producer/consumer model in order to improve real-time response. In the producer/consumer model, one originating message that is produced by one station is consumed by several stations called consumers. With Ethernet, this generates many broadcast and multicast messages that can consume the total bandwidth of the LAN. Is there another way of retaining the same physical network, but allowing separate LAN functionality? Yes there is, and it is called virtual local area networks (VLANs).

VLAN Structure

A LAN consists of stations, repeating hubs and switching hubs operating at the data link layer. LANs could be connected to other LANs if routers are used; thereby, creating an internetwork. Each LAN would then be given a network address. The best example of an internetwork is the Internet. Therefore, it is possible to have the industrial automation system on one LAN and the information system on another LAN with the two linked by a router. However, the structured wiring within the plant may not support this wiring directly.

Besides, configuring routers is more difficult than configuring VLANs. What is desired is to have the information system and industrial automation system on the same LAN, but logically separated into two LANs. That is what a VLAN can do.

Within a LAN that has all stations connected to repeating hubs, all stations hear all three types of transmissions—unicast, multicast and broadcast. In this situation, it is not possible to establish separate VLANs since there is no way of restricting traffic. A basic requirement of VLANs is the use of switching hubs. A switch learns the location of stations by observing the source MAC address present in a message received at an incoming port. The MAC address-port number association is so noted in its filtering database. All future transmissions destined to a MAC address that is stored in the switch's filtering database, will only be directed to the port associated with that MAC address unless the transmission originated on that port. If a MAC address is received with no association, the transmission is flooded to all ports (except for the received port) as if the switch were a repeating hub. The same is true for multicast and broadcast messages. Therefore, a switch provides an improvement in performance over repeating hubs by restricting unicast messages to only those stations involved, but it is this filtering capability that can be exploited for VLAN use. A single switching hub can be so configured and thus act as several independent switching hubs by creating VLAN associations to switch ports.

Port VLAN

There are several ways of creating VLANs, but the easiest to understand is the Port VLAN. Switches create an association of MAC addresses and port numbers. What needs to be added is a VLAN association. This would have to be accomplished through some configuration of a switch that can support VLANs. VLAN support is not possible with a Plug and Play switch—one with no means of altering its personality through operator intervention. For example, within a sixteen-port switch we want to create three separate VLANs numbered one to three. During configuration, we associate each port on the switch to be a VLAN. From then on, traffic within a VLAN assignment will be restricted to only those ports associated with that VLAN assignment. Using our example of three VLANs, we established VLAN1 as associated with ports 1, 2, 3 and 4. A broadcast or multicast message on port 1 would be sent only to ports 2, 3 and 4 and no others. The other VLANs would operate in a similar fashion. A unicast message would be forwarded as with any other switch. There would be a MAC address-port number association. However, added to this association would be the VLAN constraints. So if the MAC address-port number association is not present in memory for a

destination address, flooding will only occur with the VLAN port group. What happens when a destination address is specified in a transmission received on a port from another VLAN group? The transmission should be discarded.

Figure 2 shows a Port VLAN application consisting of three VLANs, although more VLANs can be added. There is only one VLAN-aware switch located in the middle of the LAN. The other switches that are not VLAN-aware are considered part of the respective VLANs. Each port on the VLAN-aware switch has an association with a common port on the switch where a server resides. This overlapping of VLANs allows any workstation in a VLAN to access the server, but workstations in separate VLANs are not known to each other.




Figure 2. In this Port VLAN application, the server in the middle is logically attached to all three VLANs.

A big advantage of Port VLAN is that it is simple to understand and use. Patch panel ports can be tagged with the associated VLAN, and it is just a simple matter of moving patch cords around to connect particular stations to particular VLANs. A simpler way of doing it is to have software do it. By reconfiguring the VLAN-aware switch, physical ports can be reassigned to different VLANs. However, what if you want to stretch your VLAN across several switches? It is possible, but you would need to have dedicated wiring for each VLAN. That is a severe restriction and, therefore, Port VLANs are best accomplished using a single VLAN-aware switch. Notice that there is no change in Ethernet frames with Port VLAN partitioning. End stations are unaware of the VLAN structure. More flexibility is gained if VLAN associations can be learned from the contents of the Ethernet frame. This is called implied tagging which allows VLANs to span multiple switches using the same cabling structure.

Frame Encoded VLAN Schemes

With Port VLAN, there is no altering of Ethernet frames or any implicit tagging within Ethernet frames. Stations are unaware of the VLAN structure. There are alternate ways of establishing VLANs if the switches being used support the various schemes. You could

simply associate particular MAC addresses to a VLAN. In this way the station assigned to the VLAN can be on any switch port and still be attached to a particular VLAN. Of course, if that station were ever replaced, all switches would need to be reconfigured for the new MAC address. Another approach to VLANs is to separate stations according to the network operating system being supported. By examining some protocol field, frames could be directed only to those stations supporting that operating system. This approach to VLANs was popular when there were several competing network operating systems with much different Ethernet frame definitions. The movement towards universal TCP/IP acceptance has now limited the frame structure choices. Another scheme is to define a proprietary protocol by coding the Ethernet frame with VLAN information. The problem with proprietary schemes is that they do not have wide industry support. To obtain wide industry support, you need an IEEE standard.

Explicit VLAN Tagging

Ethernet has been around since the mid-70s, and the maximum length frame (less preamble) was always 1518 bytes. For industrial automation, this frame size is quite large since most I/O messages are short. However, after all these years it appears that 1518 bytes are still not enough. The IEEE 802.1Q committee decided that four more bytes were needed in order to define a universally acceptable VLAN tag. There were concerns that stations and hubs could not handle an oversized frame and this new standard required a revision to IEEE 802.3. Everything we said about maximum frame size is now wrong. It is not 1518 bytes, but 1522 when VLAN tags are appended.

The IEEE 802.1Q VLAN tagging scheme is called an explicit VLAN scheme since something (VLAN tag itself) is appended to the frame versus being implied (implicit VLAN) by the contents of the frame. The four-byte tag is inserted immediately after the source address and before the Type/Length field (Fig. 3). The first two bytes are called the Tag Protocol Identifier and functions much like the Type/Length field. The contents of the two bytes are 0x8100, which is to be recognized as a VLAN tag. The following two bytes are the Tag Control Information. The remainder of the Ethernet frame stays the same except the Frame Check Sequence (FCS) must be recalculated because of the longer frame. Other than that restriction, a VLAN tag can be added or removed without affecting the contents or nature of the message.

The two-byte Tag Control Information consists of three bits for IEEE 802.1p priority levels (that has nothing to do with VLANs), one bit called the Canonical Format Indicator (CFI) and 12 bits for the VLAN identifier. With 12 bits of identifier, there could be up to 4096 VLANs. However, all ones are reserved and all zeros indicate no VLAN association, meaning

that the tag is solely used to indicate priority level. All other identifiers can be used to indicate a particular VLAN along with the 802.1p priority level of the message.

The CFI bit is used to indicate bit ordering within frames, which is an issue when communicating over non-Ethernet LANs. Since we are only interested in Ethernet LANs, the CFI bit is set to zero.

VLAN-unaware End Stations and Switches

Since 802.1Q arrived over 20 years after the invention of Ethernet, there are plenty of VLAN-unaware devices in the field. Although an end station will probably accept the elongated frame, will the software driver "choke" on receiving a 0x8100 Ethernet protocol identifier that it has never seen before? The best practice is for end stations never to see VLAN tags unless they are conditioned to do so. With the amount of legacy equipment in the field, it is a good bet the end stations are VLAN-unaware. A VLAN-aware end station is one that can receive and apply 802.1Q VLAN tags and, therefore, is termed tag-aware. However, the same is not true of switches. A VLAN-aware switch must be able to make VLAN-port associations but it may not understand 802.1Q tagging. A Port VLAN switch is a good example. A tag-aware switch understands 802.1Q tagging and can make VLAN-port associations as well.

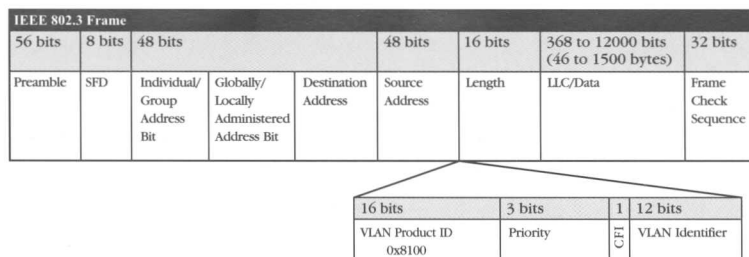


Figure 3. 802.1Q VLAN standard inserts a four-byte tag into a standard Ethernet frame.

VLAN Edge Switches

If a VLAN-aware station initiated a transmission received on a port of a tag-aware switch, it is a simple matter to read the value of the VLAN assignment and forward the frame intact to those ports in its filtering database for that particular VLAN assignment. However, if a transmission is instead received from a VLAN-unaware station, the tag-aware switch must append a VLAN tag equivalent to the VLAN association established previously for the received frame. This association could be based on the MAC address, protocol ID or port location as discussed earlier. Whatever the association rule was for the VLAN, the identifier for that VLAN must be the same as applied to the VLAN tag and the new frame forwarded to the output port or ports indicated in the switch's filtering database.

In order to limit VLAN tags from being propagated to VLAN-unaware end stations, the tag-aware switch

must have the capability of removing VLAN tags at output ports. This capability can be found in an edge switch that resides on the boundary of VLAN-aware and VLAN-unaware domains. An edge switch can read a VLAN tag from a VLAN-aware station or append a VLAN tag to a frame from a VLAN-unaware station and take appropriate forwarding action. Before it forwards the frame to one of its output ports, it looks in its table if the VLAN tag is to remain or be removed. If the message is going to VLAN-unaware stations, then the VLAN tag should be stripped. If it is going on to core VLAN switches, then it should be retained.

VLAN Core Switches

Core switches understand VLAN tags and reside in the backbone of the LAN and are usually only connected to edge switches. Therefore, their forwarding rules are much simpler and faster to implement. All incoming frames will have VLAN tags and all outbound frames will retain these tags. The filtering database could consist of only the 4094 possible VLANs and output port assignments. No source addressing would need to be learned. In actuality, an edge switch could be configured as a core switch, and since it would probably be too confusing to have two types of VLAN-aware switches in the plant, restricting use to only edge switches could be the answer. Even though 4094 VLANs are possible according to the 802.1Q standard, not all switches can support that many VLANs simultaneously. Could you imagine the complexity of configuring and maintaining this many VLANs?

Mobility

It would be convenient to be able to take your laptop and connect it to any available spare port on a switch within the LAN and examine the operation of an industrial automation system on a particular VLAN. In order to effectively achieve this functionality, the laptop should be VLAN-aware and the attached switch must be programmed to allow access for that particular VLAN by having a valid VLAN-port association that would reach the VLAN desired. Using a VLAN-unaware laptop with implicit tagging would make the task even more difficult, but not impossible. Reconfiguration of the various switches in the path of the VLAN may be required in order to open up the port attached to the laptop. The use of Port VLANs would be impractical.




Figure 4. The most flexible VLAN arrangement can be achieved by the use of 802.1Q tags. Edge switches allow the use of both VLAN-aware and VLAN-unaware end stations.

Figure 4 shows a typical LAN incorporating 802.1Q tagging with edge switches each connected to one core switch using a single cable. Within the VLAN-aware domain, edge switches must transmit VLAN-tagged frames to identify frame-VLAN associations. For any edge switch to have access to all possible VLANs (to ensure mobility), the port connected to the core switch must be associated with all possible VLANs.

CONCLUSION

VLANs are an effective means of portioning a larger LAN into manageable subsets. VLANs restrict the broadcast domain, improve performance and security, and they are ideal for isolating industrial automation systems from IT systems while retaining the plant's structural wiring. The simplest of VLANs to implement are Port VLANs, but the most effective VLAN scheme is the IEEE 802.1Q VLAN tagging standard that improves mobility by allowing a user to potentially access any VLAN from any point on the LAN.

REFERENCES

- The Switch Book*, Rich Seifert, 2000, Wiley Computer Publishing
- Ethernet The Definite Guide*, Charles E. Spurgeon, 2000, O'Reilly & Associates, Inc.
- International Standard ISO/IEC 8802-3 ANSI/IEEE Std 802.3*, 2000, The Institute of Electrical and Electronics Engineers, Inc.
- Commercial Building Telecommunications Cabling Standard*, TIA/EIA-568-A, 1995, Telecommunications Industry Association
- Virtual Bridged Local Area Networks IEEE Std 802.1Q™*, 2003 Edition, The Institute of Electrical and Electronics Engineers, Inc.

CONTEMPORARY CONTROLS
www.ccontrols.com
www.ctrlink.com

Past issues of the Extension are available. If you would like a copy, please send your request to info@ccontrols.com